Applied and Computational Harmonic Analysis 10, 234–253 (2001)
doi:10.1006/acha.2000.0343, available online at http://www.idealibrary.com on IDEAL®

Complex Wavelets for Shift Invariant Analysis and Filtering of Signals

Nick Kingsbury

Department of Engineering, Signal Processing Group, University of Cambridge, Cambridge CB2 1PZ, United Kingdom
E-mail: ngk@eng.cam.ac.uk

This paper describes a form of discrete wavelet transform, which generates *complex* coefficients by using a *dual tree* of wavelet filters to obtain their real and imaginary parts. This introduces limited redundancy (2^m : 1 for m-dimensional signals) and allows the transform to provide approximate shift invariance and directionally selective filters (properties lacking in the traditional wavelet transform) while preserving the usual properties of perfect reconstruction and computational efficiency with good well-balanced frequency responses. Here we analyze why the new transform can be designed to be shift invariant and describe how to estimate the accuracy of this approximation and design suitable filters to achieve this. We discuss two different variants of the new transform, based on odd/even and quarter-sample shift (Q-shift) filters, respectively. We then describe briefly how the dual tree may be extended for images and other multi-dimensional signals, and finally summarize a range of applications of the transform that take advantage of its unique properties.

1. INTRODUCTION

Signal compression (coding) has for some time been a very active area for signal processing research, and the wavelet transform has established an impressive reputation as a tool for this, especially for images and motion video. Many researchers have also tried to use wavelets for signal analysis and reconstruction but the results have tended to be disappointing. Here we consider possible reasons for this and propose the dual-tree *complex* wavelet transform as a useful tool for overcoming some of these problems.

The discrete wavelet transform (DWT) is most commonly used in its maximally decimated form (Mallat's dyadic filter tree [1, 2]). This works well for compression but its use for other signal analysis and reconstruction tasks has been hampered by two main disadvantages:

- Lack of *shift invariance*, which means that small shifts in the input signal can cause major variations in the distribution of energy between DWT coefficients at different scales.
- Poor *directional selectivity* for diagonal features, because the wavelet filters are separable and real.

1063-5203/01 $35.00

A well-known way of providing shift invariance is to use the undecimated form of the dyadic filter tree, which is implemented most efficiently by the *algorithme à trous* [2, Sect. 5.5.2]. However, this still suffers from substantially increased computation requirements compared to the fully decimated DWT and also exhibits high redundancy in the output information, making subsequent processing expensive too.

In [10, 11], we introduced a more computationally efficient approach to shift invariance, the dual-tree complex wavelet transform (DT CWT). Furthermore the DT CWT also gives much better directional selectivity when filtering multi-dimensional signals. In summary, it has the following properties:

- Approximate **shift invariance**;
- Good **directional selectivity** in 2-dimensions (2-D) with Gabor-like filters (also true for higher dimensionality, m-D);
- **Perfect reconstruction** (PR) using short linear-phase filters;
- **Limited redundancy**, independent of the number of scales, 2 : 1 for 1-D (2^m : 1 for m-D);
- **Efficient order-N computation**—only twice the simple DWT for 1-D (2^m times for m-D).

Another approach both to shift invariance and to directional selectivity was proposed by Simoncelli *et al.* [4] and was based on Laplacian pyramids and steerable filters, designed in the frequency domain. Laplacian pyramids were first proposed by Burt and Adelson [3] and, in that form, also produced limited redundancy of 2 : 1 for 1-D and 1.33 : 1 for 2-D signals. In addition, they could provide moderate levels of shift invariance if the lowpass filters of the pyramid were designed to be sufficiently narrow in bandwidth. However, only with Simoncelli's steerable filters was directional selectivity achieved in 2-D, and then the redundancy increased to 5.33 : 1.

An alternative approach to directional selectivity was pioneered by Bamberger and Smith [5], using a combination of 2-D resampling methods and separable polyphase filter bank techniques. This produced fan-shaped directional decompositions of the input signal with maximal decimation (no redundancy). However, due to being maximally decimated, such filters would be likely to exhibit high shift dependency. Furthermore, it is unclear how to obtain efficient multiscale decompositions from such a filter bank.

The complex wavelet methods, presented here, are believed to offer a useful combination of properties not available from these earlier approaches—principally perfect reconstruction, greater directional selectivity, and a natural multiscale decomposition. In [12] we proposed a way of analyzing the shift invariant properties of the DT CWT and here we expand on these basic ideas and develop them somewhat further.

2. THE DUAL FILTER TREE

Our work with complex wavelets for motion estimation [7] showed that complex wavelets could provide approximate shift invariance. Unfortunately we were unable to obtain PR and good frequency characteristics using short support complex FIR filters in a single tree (e.g., Fig. 1, tree a).

However, we observed that we can also achieve approximate shift invariance with a *real* DWT, by doubling the sampling rate at each level of the tree. For this to work, the samples

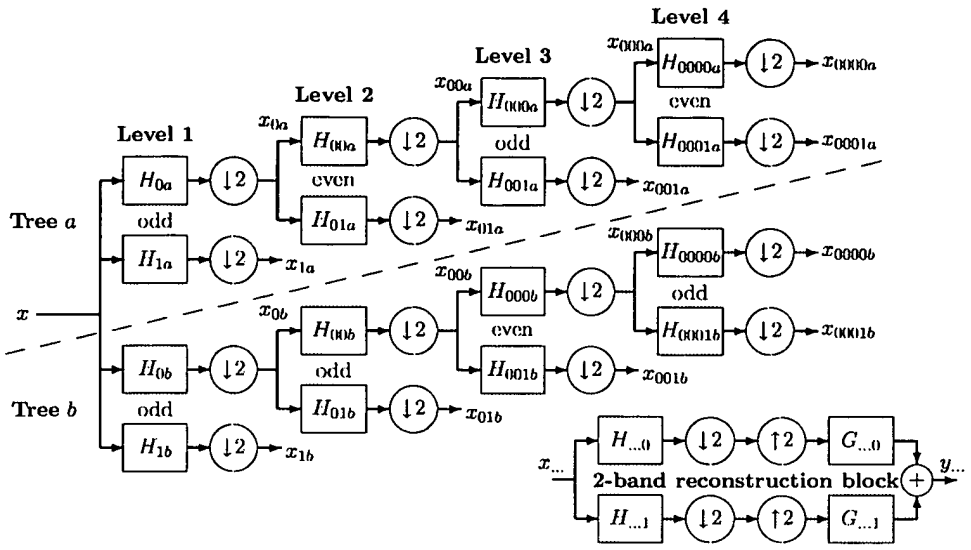

FIG. 1. The dual-tree complex wavelet transform (DT CWT), comprising two trees of real filters, a and b, which produce the real and imaginary parts of the complex coefficients. Odd and even length biorthogonal linear-phase filters are placed as shown to achieve the correct relative signal delays.

must be evenly spaced. One way to double all the sampling rates in a conventional wavelet tree, such as tree a of Fig. 1, is to eliminate the down-sampling by 2 after the level 1 filters, H_{0a} and H_{1a}. This is equivalent to having two parallel fully decimated trees, a and b in Fig. 1, provided that the delays of filters H_{0b} and H_{1b} are one sample offset from the delays of H_{0a} and H_{1a}, which ensures that the level 1 downsamplers in tree b pick the opposite samples to those in tree a.

While this approach works at level 1, we find that *below* level 1 the tree b samples do not interpolate midway between the tree a samples, because of the lower sampling rates at these levels. Hence shift invariance would not be achieved below level 1. To achieve this at every level, the total delay difference for a given level and all previous levels must sum to one sample period at the input sample rate of the given level. Hence the filters below level 1 in one tree must provide delays that are *half a sample* different (at each filter's input rate) from those in the opposite tree. For linear phase filters, this requires *odd-length* filters in one tree and *even-length* filters in the other.

Greater symmetry between the two trees occurs if each tree uses odd and even filters alternately from level to level, but this is not essential. In Fig. 2a we show the positions of the wavelet basis functions when the filters are arranged to be odd and even as in Fig. 1. Note the vertical alignment of these bases at each scale, such that the tree b scaling functions interpolate midway between those of tree a, while the tree b wavelets are aligned with those of tree a but with a quadrature phase shift in the underlying oscillation.

To invert the DT CWT, each tree in Fig. 1 is inverted separately using biorthogonal filters $G_{...}$, designed for perfect reconstruction with the corresponding analysis filters $H_{...}$ in the 2-band reconstruction block, shown lower right. Finally the two tree outputs are averaged in order to obtain an approximately shift invariant system. This system

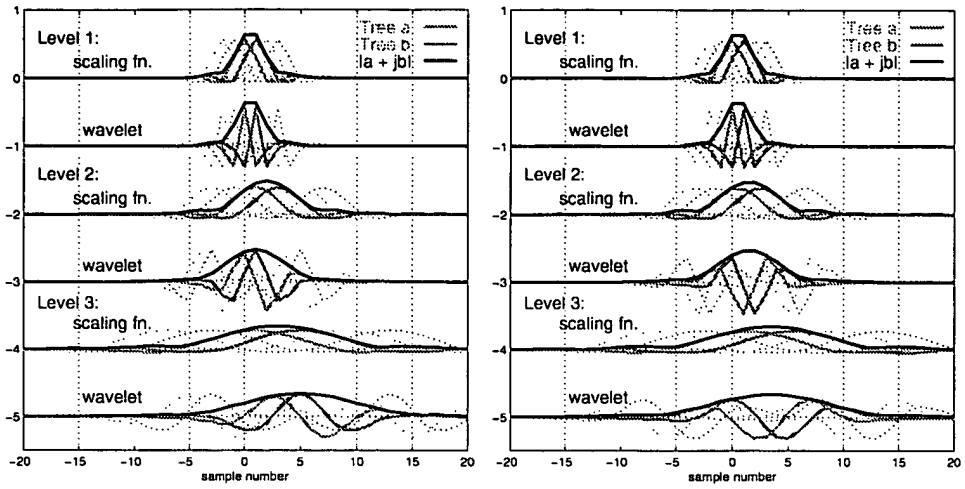

FIG. 2. (a) Basis functions (reconstruction impulse responses) of the odd and even filters in Fig. 1 for levels 1 to 3. Tree a bases are shown in light gray and tree b in darker gray. The magnitudes of the complex bases, formed by combining the two trees, are shown in black. The bases for adjacent sampling points are shown dotted. (b) Equivalent bases for the Q-shift tree of Fig. 3.

is a wavelet *frame* [2] with redundancy two; and if the filters are designed such that the analysis and reconstruction filters have very similar frequency responses (i.e., are almost orthogonal, as is the case for the filters given later in Table 1), then it is an almost tight frame, which means that energy is approximately preserved when signals are transformed into the DT CWT domain. The basis functions in Fig. 2a were obtained by injecting unit pulses separately into the inverse DT CWT at each scale in turn. The real and imaginary parts were obtained by injecting the unit pulses into trees a and b in turn.

3. THE Q-SHIFT DUAL-TREE FILTERS

Unfortunately there are certain problems with the odd/even filter approach:

- The sub-sampling structure is not very symmetrical (Fig. 2a shows that the magnitudes of the wavelets at levels 2 and 3 are not aligned with the corresponding scaling functions);
- The two trees have slightly different frequency responses;
- The filter sets must be biorthogonal, rather than orthogonal, because they are linear phase. This means that energy is not preserved between the signal and transform domains.

The latter two problems are not severe for most purposes and their effects can be minimized with relatively long filters (13 to 19 taps), but there is then a computation penalty. The first problem is more fundamental and has implications for subsequent hierarchical algorithms, such as hidden Markov trees [16, 18].

To overcome all of the above (and with no significant penalties), we now propose a *Q-shift* dual tree, as in Fig. 3, in which all the filters beyond level 1 are even length,

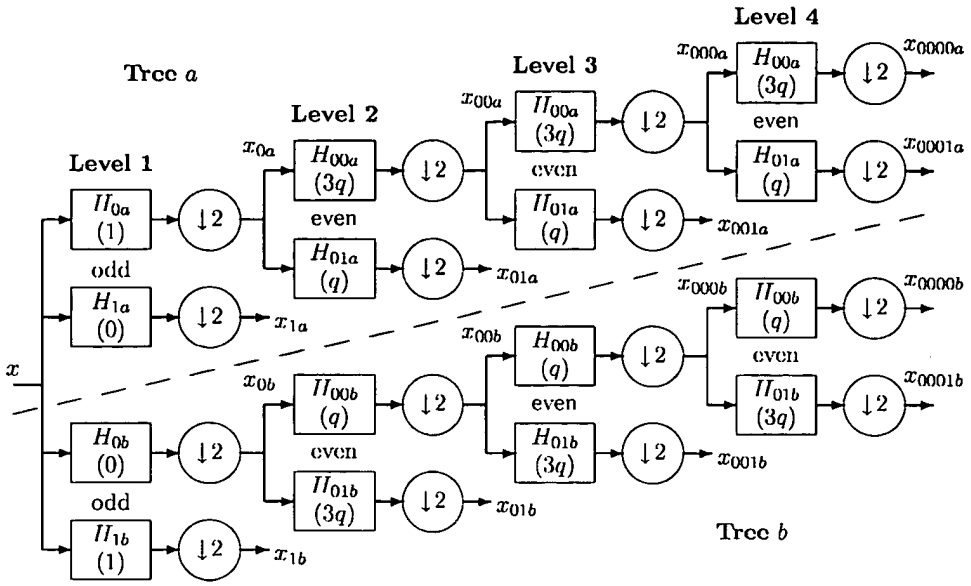

FIG. 3. The Q-shift version of the DT CWT, giving real and imaginary parts of complex coefficients from tree a and tree b, respectively. Figures in brackets indicate the delay for each filter, where $q = \frac{1}{4}$ sample period.

but they are no longer strictly linear phase. Instead they are designed to have a group delay of approximately $\frac{1}{4}$ sample ($+q$). The required delay difference of $\frac{1}{2}$ sample ($2q$) is then achieved by using the time reverse of the tree a filters in tree b so that the delay then becomes $3q$ (assuming that all length-$2n$ filters have coefficients from z^{n-1} to z^{-n}). Furthermore, because the filter coefficients are no longer symmetric, it is now possible to design the perfect-reconstruction filter sets to be orthonormal (like Daubechies filters), so that the reconstruction filters are just the time reverse of the equivalent analysis filters in both trees. Hence *all* filters beyond level 1 are derived from the *same* orthonormal prototype set. The design of Q-shift filters is discussed later in Section 6.

The improved sampling symmetry of the Q-shift filters is shown by the basis functions in Fig. 2b. Note that for the Q-shift CWT the magnitude of each complex wavelet basis is aligned with the corresponding complex scaling function basis, and that each of these is centerd between a pair of adjacent complex bases from the previous (finer) level. In this way, each complex wavelet coefficient at level k has two complex *children* located symmetrically above it at level $k-1$. For the odd/even DT CWT in Fig. 2a, such alignments do not occur.

4. SHIFT INVARIANCE

In order to examine the shift invariant properties of the dual tree in either the odd/even or Q-shift forms, consider what happens when we choose to retain the coefficients of just one type (wavelet or scaling function) from just one level of the dual tree. For example, we might choose to retain only the level-3 wavelet coefficients x_{001a} and x_{001b}, and set all others to zero. If the signal y, reconstructed from just these coefficients, is free of aliasing then we define the transform to be shift invariant at that level. This is because absence of

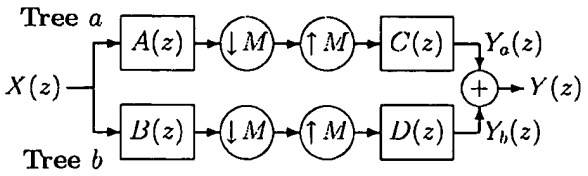

FIG. 4. Basic configuration of the dual tree if either wavelet or scaling-function coefficients from just level m are retained ($M = 2^m$).

aliasing implies that a given subband is a linear time-invariant (LTI) system and therefore is fully characterized by a single impulse response and a corresponding z-transfer function. In this context we define a subband as comprising all coefficients from *both* trees at a given level and of a given type (either wavelet or scaling function).

Figure 4 shows the simplified analysis and reconstruction parts of the dual tree when coefficients of just one type and level are retained. All down-sampling and up-sampling operations are moved to the outputs of the analysis filter banks and the inputs of the reconstruction filter banks, respectively, and the cascaded filter transfer functions are combined. $M = 2^m$ is the total down/up-sampling factor. For example, if x_{001a} and x_{001b} from fig. 1 are the only retained coefficients, then the sub-sampling factor $M = 8$, and $A(z) = H_{0a}(z)\, H_{00a}(z^2)\, H_{001a}(z^4)$, the transfer function from x to x_{001a}. The transfer function $B(z)$ (from x to x_{001b}) is obtained similarly using $H_{\ldots b}(z)$; as are the inverse functions $C(z)$ and $D(z)$ from $G_{\ldots a}(z)$ and $G_{\ldots b}(z)$, respectively.

It is a standard result of multi-rate analysis that a signal $U(z)$, which is down-sampled by M and then up-sampled by the same factor (by insertion of zeros), becomes $(1/M)\sum_{k=0}^{M-1} U(W^k z)$, where $W = e^{j2\pi/M}$. Applying this result to Fig. 4 gives

$$Y(z) = Y_a(z) + Y_b(z) = \frac{1}{M}\sum_{k=0}^{M-1} X(W^k z)[A(W^k z)\, C(z) + B(W^k z)\, D(z)]. \qquad (1)$$

The aliasing terms in this summation correspond to those for which $k \neq 0$, because only the term in $X(z)$ (when $k = 0$ and $W^k = 1$) corresponds to a linear time (shift) invariant response. For shift invariance, the aliasing terms must be negligible, so we must design $A(W^k z)\, C(z)$ and $B(W^k z)\, D(z)$ either to be very small or to cancel each other when $k \neq 0$. Now W^k introduces a frequency shift equal to kf_s/M to the filters A and B (f_s is the input sampling frequency), so for larger values of k the shifted and unshifted filters have negligible passband overlap and it is quite easy to design the functions $B(W^k z)\, D(z)$ and $A(W^k z)\, C(z)$ to be very small over all frequencies, $z = e^{j\theta}$. But at small values of k (especially $k = \pm 1$) this becomes virtually impossible due to the significant transition band widths of short-support filters. Here it is necessary to design for cancellation when the two trees are combined, and separate strategies are required depending on whether the filters are lowpass (for scaling functions) or bandpass (for wavelets).

First consider the scaling function (lowpass) filters. At level m in the dual tree, the lowpass filters have passbands $\{-f_s/2M \leftrightarrow f_s/2M\}$. The W^k terms in Eq. (1) shift the passbands in multiples of f_s/M. If $A(z)$ and $C(z)$ have similar frequency responses (as required for near-orthogonal filter sets) and significant transition band widths, it is not possible to make $A(W^{\pm 1} z)\, C(z)$ small at all frequencies $z = e^{j\theta}$, because the transition

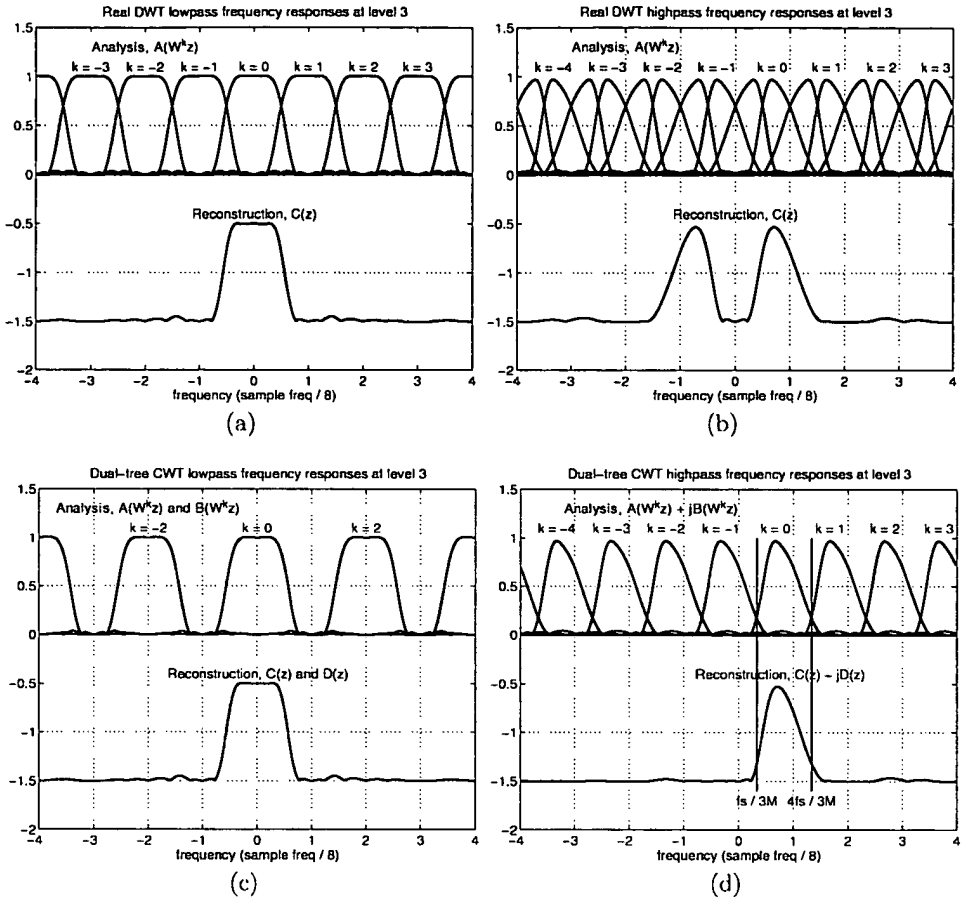

FIG. 5. Frequency responses of analysis and reconstruction filters (from Table 1) at level 3, showing aliasing terms. Plots (a) and (b) show the responses of the lowpass and bandpass filters in a single wavelet tree, while (c) and (d) show the equivalent responses for the dual tree and demonstrate the reduced overlap of the reconstruction filters with the frequency-shifted (aliased) analysis filters. The horizontal axes are in units of f_s/M where $M = 8$ for level 3. The reconstruction responses are offset vertically by -1.5 to avoid confusion.

bands of the shifted analysis filter $A(W^{\pm 1}z)$ overlap with those of the reconstruction filter $C(z)$ as shown in Fig. 5a. However, we can quite easily make $A(W^{\pm 2}z)\,C(z)$ small since the frequency shift is twice as great. Hence for the lowpass case, we design $B(W^k z)\,D(z)$ to cancel $A(W^k z)\,C(z)$ when k is odd by letting

$$B(z) = z^{\pm M/2}A(z) \qquad \text{and} \qquad D(z) = z^{\mp M/2}C(z) \tag{2}$$

so that $B(W^k z)\,D(z) = (-1)^k A(W^k z)\,C(z)$. The effect of this cancellation is shown in Fig. 5c and is equivalent to halving the down-sampling factor M.

Now consider the wavelet (bandpass) filters. Here we find that the edges of the negative frequency (lower) passband of C, covering the range $\{-f_s/2M \leftrightarrow -f_s/M\}$ as shown in Fig. 5b, will tend to overlap with the edges of the positive frequency (upper) passband of A, that gets shifted either to $\{0 \leftrightarrow -f_s/2M\}$ or to $\{-f_s/M \leftrightarrow -3f_s/2M\}$ when $k = -1$ or

-2, respectively. Similarly the upper passband of C will overlap the lower passband of A when $k = +1$ or $+2$. The filters D and B, which are designed to have almost the same responses as C and A, will behave in the same way. Since the aliasing terms are always caused by the overlap of opposing-frequency passbands, whereas the wanted terms ($k = 0$) are produced by overlap of same-frequency passbands, the solution here is to give B and D upper and lower passbands of opposite polarity while A and C have passbands of the same polarity (or vice versa). The need to discriminate between positive and negative frequency components in this way suggests the use of *complex* filters!

Suppose that we have prototype complex filters $P(z)$ and $Q(z)$, each with just a single passband $\{f_s/2M \leftrightarrow f_s/M\}$ and negligible gain at all negative frequencies, and that we let

$$\begin{aligned} A(z) &= 2\,\Re[P(z)] = P(z) + P^*(z) \\ B(z) &= 2\,\Im[P(z)] = -j[P(z) - P^*(z)] \\ C(z) &= 2\,\Re[Q(z)] = Q(z) + Q^*(z) \\ D(z) &= -2\,\Im[Q(z)] = j[Q(z) - Q^*(z)], \end{aligned} \tag{3}$$

where $\Re[\ ]$ and $\Im[\ ]$ take real and imaginary parts, and conjugation is given by $P^*(z) = \sum_r p_r^* z^{-r}$. The process of conjugating a filter that has a positive frequency passband produces one with an equivalent negative frequency passband. Hence P and Q are filters corresponding to the upper passbands of A and C while P^* and Q^* correspond to their lower passbands. This also applies to filters B and D, except that their lower passbands are negated. Additionally we may observe that the (real) impulse responses of B and D are the Hilbert transforms of those of A and C.

Returning to Fig. 5b, the main overlap terms are due to opposing frequency passbands of the form $Q^*(z)\,P(W^k z)$ for $k = -1, -2$ and $Q(z)\,P^*(W^k z)$ for $k = 1, 2$. These cancel when $B(W^k z)\,D(z)$ is added to $A(W^k z)\,C(z)$ in (1), because, for all k,

$$\begin{aligned} A(W^k z)\,C(z) + B(W^k z)\,D(z) &= [P(W^k z) + P^*(W^k z)]\,[Q(z) + Q^*(z)] \\ &\quad + (-j)[P(W^k z) - P^*(W^k z)]\; j[Q(z) - Q^*(z)] \\ &= 2P(W^k z)\,Q(z) + 2P^*(W^k z)\,Q^*(z). \end{aligned} \tag{4}$$

Hence we now need only design the filters such that the positive frequency complex filter $Q(z)$ does not overlap with shifted versions of the similar filter $P(z)$. This is quite easy since the filter bandwidths are only $f_s/2M$ while the shifts are in multiples of f_s/M. For octave band filters in which the upper transition band is twice as wide as the lower transition band, this is satisfied if the pass and transition bands lie within the frequency range $\{f_s/3M \leftrightarrow 4f_s/3M\}$. Figure 5d shows the frequency responses of $P(W^k z)$ and $Q(z)$ from the right side of Eq. (4) and demonstrates that negligible overlap of the responses occurs for $k \neq 0$, as long as the pass and transition bands of P and Q are constrained to lie within the two vertical lines at $f_s/3M$ and $4f_s/3M$. The second term on the right side of Eq. (4) involves $P^*(W^k z)$ and $Q^*(z)$, which will just give the mirror image of the responses in Fig. 5d, and hence the same amount of overlap.

The formulations in Eqs. (3) show that the bandpass filter responses for trees a and b (A, B for analysis; C, D for reconstruction) should be regarded as the *real* and *imaginary* parts of complex responses (P for analysis; Q for reconstruction) that have passbands only

on one side of zero frequency. This is the key justification for using complex wavelets to achieve shift invariance for the wavelet coefficients.

It turns out that we may regard the pairs of scaling function coefficients from the two trees as real and imaginary parts too, although there is the alternative option of simply regarding the scaling function coefficients from tree b as interpolating mid-way between the corresponding ones from tree a, which for some applications turns out to be more useful and is the most natural interpretation of Eqs. (2).

In practice, filters with compact support will not have zero gain in their stop bands and the aliasing terms in Eq. (1) will not be zero. Furthermore for the odd/even filters the cancellation of other unwanted terms, when the two trees are combined, will not be exact because the odd-length filters cannot have precisely the same frequency responses as the even-length ones. So a typical DT CWT will only be *approximately* shift invariant. However, we shall show that good performance is possible with quite low complexity filters.

A useful way of quantifying the shift dependence of a transform is to examine Eq. (1) and determine the ratio of the total energy of the unwanted aliasing transfer functions (the terms with $k \neq 0$) to the energy of the wanted transfer function (when $k = 0$), as given by

$$R_a = \frac{\sum_{k=1}^{M-1} \mathcal{E}\{A(W^k z)\, C(z) + B(W^k z)\, D(z)\}}{\mathcal{E}\{A(z)\, C(z) + B(z)\, D(z)\}}, \tag{5}$$

where $\mathcal{E}\{U(z)\}$ calculates the energy, $\sum_r |u_r|^2$, of the impulse response of a z-transfer function, $U(z) = \sum_r u_r z^{-r}$. Of course, $\mathcal{E}\{U(z)\}$ may also be calculated as the integral of the squared magnitude of the frequency response, $(1/2\pi) \int_{-\pi}^{\pi} |U(e^{j\theta})|^2\, d\theta$ from Parseval's theorem, which links more naturally with the interpretation in Fig. 5.

5. ODD/EVEN FILTER DESIGN METHOD

We first suggest a way to design filters for the odd/even DT CWT which achieve good shift invariance.

For the lowpass filters, Eq. (2) implies that the tree b samples should interpolate midway between the tree a samples, effectively doubling the sampling rate, as shown in Fig. 2. This may be achieved by two identical odd-length lowpass filters at level 1, offset by 1 sample delay, and then by pairs of odd and even length filters at further levels to achieve an extra delay difference of $M/4$ samples, so as to make the total delay difference $M/2$ samples at each level m, where $M = 2^m$.

The responses of $A(z)$ and $B(z)$ also need to match, which can only be achieved approximately for odd/even filters beyond level 1. We do this by designing the even-length filter H_{00a} to give minimum mean squared error in the approximation

$$z^{-2}\, H_{0a}(z)\, H_{00a}(z^2) \approx H_{0b}(z)\, H_{00b}(z^2), \tag{6}$$

where H_{00b} is assumed to be the same odd-length design as the two level 1 filters, such that $H_{00b}(z) = H_{0b}(z) = z^{-1} H_{0a}(z)$. In this case solving for H_{00a} is just a matrix pseudo-inverse problem.

Then the highpass filter H_{01a} can be designed to form a perfect reconstruction set with H_{00a} such that the reconstruction filters G_{00a} and G_{00b} also match each other closely.

TABLE 1
Coefficients of the (13,19)-Tap and (12,16)-Tap Odd/Even Filters Used in the Results of This Paper

Odd $H_{...0}$ 13-tap	Odd $H_{...1}$ 19-tap	Even $H_{...0}$ 12-tap	Even $H_{...1}$ 16-tap
	−0.0000706		
	0		−0.0004645
−0.0017581	0.0013419		0.0013349
0	−0.0018834	−0.0058109	0.0022006
0.0222656	−0.0071568	0.0166977	−0.0130127
−0.0468750	0.0238560	−0.0000641	0.0015360
−0.0482422	0.0556431	−0.0834914	0.0869008
0.2968750	−0.0516881	0.0919537	0.0833552
0.5554688	−0.2997576	0.4807151	−0.4885957
0.2968750	0.5594308	0.4807151	0.4885957
−0.0482422	−0.2997576	0.0919537	−0.0833552
⋮	⋮	⋮	⋮

Finally the symmetry of the odd-length highpass filters and the anti-symmetry of the even-length highpass filters produce the required phase relationships between the positive and negative frequency passbands, and Eqs. (3) are approximately satisfied too. These odd and even length filters can then be used for all subsequent levels of the transform, in accordance with Fig. 1.

Good shift invariance (and wavelet smoothness) requires that frequency response sidelobes of the cascaded multirate filters should be small. This is achieved if each lowpass filter has a stopband covering $\frac{1}{3}$ to $\frac{2}{3}$ of its sample rate, so as to reject the image frequencies due to subsampling in the next lowpass stage. If the highpass filters then mirror this characteristic, the conditions for no overlap of the shifted bandpass responses in Eq. (4) are automatically satisfied.

To illustrate this process, we have designed two linear-phase PR biorthogonal filter sets which meet the above conditions quite well and are also nearly orthogonal. For the odd-length set, we designed (13,19)-tap filters using the (1-D) transformation of variables method [8]. Then for the even-length set, we designed a (12,16)-tap even-length set to minimize the error in Eq. (6). The analysis filter coefficients are listed in Table 1 (the reconstruction filters are obtained by negating alternate coefficients and swapping bands). These filters may be implemented efficiently using ladder structures, the odd filter pair requiring 4 multiplies and 6 additions per input sample, and the even pair 7.5 multiplies and 7 additions. Figure 6 shows the frequency responses of a 4-level reconstruction filter bank when the coefficients from the two trees are combined to form complex coefficients (the scaling function coefficients are also combined in this way). The analysis filters are very similar. Note the absence of gain at negative frequencies. We have implemented these filters and have found them to be good for many applications. However, other options do exist, as any combination of odd-length and even-length biorthogonal linear phase filters could in theory be used, although with varying levels of shift invariance and wavelet smoothness.

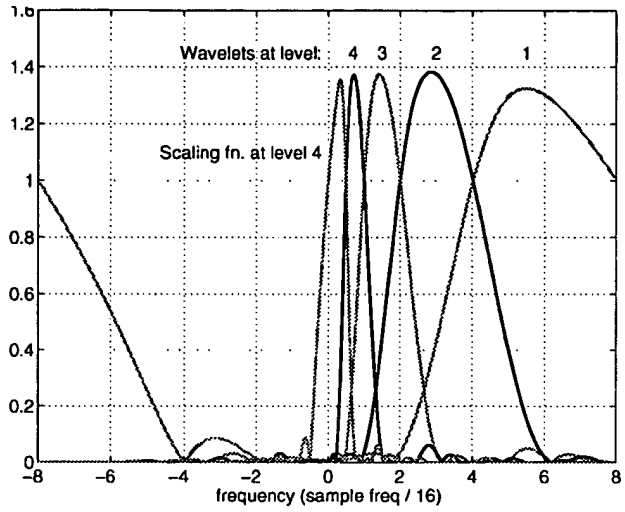

FIG. 6. Frequency responses of complex wavelets at levels 1 to 4 and of the level 4 scaling function.

6. Q-SHIFT FILTER DESIGN METHOD

The key to designing filters for the Q-shift version of the DT CWT lies in finding a good even-length lowpass filter with a delay of $\frac{1}{4}$ sample which also satisfies the standard orthonormal perfect reconstruction condition of two-band filter banks [6].

To achieve a lowpass filter $H_L(z)$ of length $2n$ with a delay that approximates $\frac{1}{4}$ sample, we find that the simplest approach is to design a linear-phase lowpass FIR filter $H_{L2}(z)$ of length $4n$ with half of the desired bandwidth and twice the desired delay, and then to select alternate filter coefficients to obtain H_L. Hence, if

$$H_{L2}(z) = H_L(z^2) + z^{-1}\, H_L(z^{-2}), \tag{7}$$

where $H_L(z)$ contains coefficients from z^{n-1} to z^{-n}, then H_{L2} automatically is linear phase with a delay of $\frac{1}{2}$ sample (see Fig. 7). As long as $H_{L2}(z)$ has very little gain above a quarter of its sampling frequency, there will be negligible aliasing caused by the subsampling and $H_L(z)$ will have a delay closely approximating $\frac{1}{4}$ sample.

We then adjust H_L such that the squared gain of H_{L2} in a suitably chosen stopband is minimized subject to H_L satisfying the usual perfect reconstruction (PR) condition

$$H_L(z)\; H_L(z^{-1}) + H_L(-z^{-1})\; H_L(-z) = 2. \tag{8}$$

To achieve this, H_L may be derived from a polyphase matrix [6], factorized into a cascade of orthonormal rotations $\mathbf{R}(\theta_i)$ and delays $\mathbf{Z}$, such that

$$\begin{bmatrix} H_L(z) \\ z^{-1} H_L(-z^{-1}) \end{bmatrix} = \mathbf{R}(\theta_n)\, \mathbf{Z}\, \mathbf{R}(\theta_{n-1})\, \mathbf{Z}\, \ldots\, \mathbf{R}(\theta_1) \begin{bmatrix} 1 \\ z^{-1} \end{bmatrix}, \tag{9}$$

where

$$\mathbf{R}(\theta_i) = \begin{bmatrix} \cos\theta_i & \sin\theta_i \\ -\sin\theta_i & \cos\theta_i \end{bmatrix} \quad \text{and} \quad \mathbf{Z} = \begin{bmatrix} z & 0 \\ 0 & z^{-1} \end{bmatrix}.$$

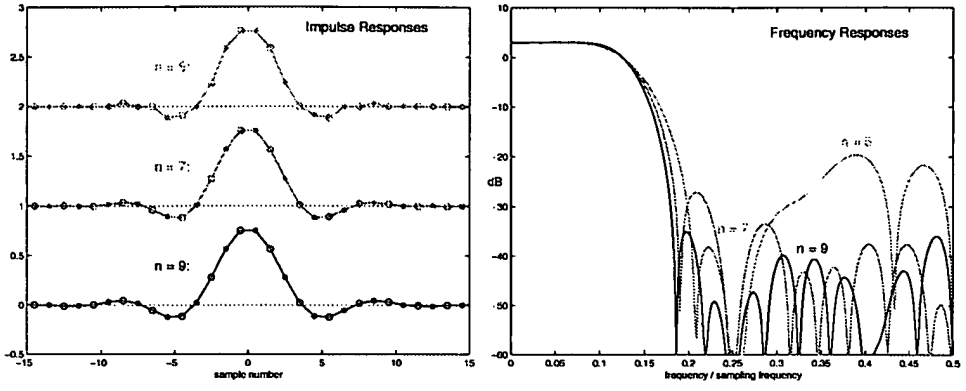

FIG. 7. Impulse and frequency responses of the oversampled filter $H_{L2}(z)$ for $n = 5$, 7, and 9. The samples of $H_L(z^2)$ are shown as circles and those of $z^{-1}H_L(z^{-2})$ as asterisks.

If $H_L(z)$ is to have at least one zero at $z = -1$, the n rotations θ_i must sum to $\pi/4$ (plus $k\pi$). (To show this, set $z = 1$ and require that $H_L(-z^{-1}) = 0$ in Eq. (9).) This leaves only $n - 1$ angles to optimize, instead of $2n$ coefficients of H_L, and it automatically satisfies the PR condition. There are several standard methods, such as in [6], which can be used to optimize $\theta_1 \ldots \theta_{n-1}$. We have found that a good criterion for the optimization is to minimize the total energy of the frequency response $H_{L2}(e^{j\omega})$ over the range $0.36\pi \leq \omega \leq \pi$. (The lower limit of 0.36π is found by experiment to give an adequate filter transition bandwidth above and below $\pi/4$.)

An interesting case occurs for $n = 5$, because a good solution exists in which only 6 of the 10 coefficients of H_L are non-zero, yielding the following 6-tap filter:

$$H_L(z) = 0.03516384z^4 - 0.08832942z^2 + 0.23389032z + 0.76027237 + 0.58751830z^{-1} - 0.11430184z^{-3}.$$

This is obtained from $\theta = \{0,\ 1.81,\ 0.81,\ -1.62,\ 0\}\pi/4$.

Smoother wavelets and scaling functions may be obtained if n is increased. For $n = 5$, 7, and 9, coefficients of H_L that we have obtained are listed in Table 2. The filters are normalized for unit energy and the dc gain is $\sqrt{2}$. Even values of n are also valid, but seem to offer slightly worse tradeoffs.

For all of the above cases, the filters after the first level in Fig. 3 are given by $H_{00a}(z) = z^{-1}H_L(z^{-1})$, $H_{01a}(z) = H_L(-z)$, $H_{00b}(z) = H_L(z)$, and $H_{01b}(z) = z^{-1}H_L(-z^{-1})$. The reconstruction filters are just the time reverses of these (i.e., trees a and b are swapped).

Figure 7 shows the impulse and frequency responses of the length-$4n$ linear-phase filters H_{L2} which result from the above design process. It is clear that the larger values of n produce smoother filters with significantly lower sidelobes in the stopband ($f \geq 0.18$ or $\omega \geq 0.36\pi$) and with somewhat narrower transition bands.

7. SHIFT INVARIANT PERFORMANCE

Having designed a range of filters for both the odd/even and Q-shift forms of the DT CWT, it is interesting to investigate their degrees of shift invariance, based on the aliasing energy ratio R_a, defined in Eq. (5) for the equivalent system in Fig. 4.

TABLE 2
Filter Coefficients of $H_L(z)$ for $n = 5, 7$, and 9

6-tap H_L	14-tap H_L	18-tap H_L
		−0.00228413
		0.00120989
	0.00325314	−0.01183479
	−0.00388321	0.00128346
0.03516384	0.03466035	0.04436522
0	−0.03887280	−0.05327611
−0.08832942	−0.11720389	−0.11330589
0.23389032	0.27529538	0.28090286
0.76027237	0.75614564	0.75281604
0.58751830	0.56881042	0.56580807
0	0.01186609	0.02455015
−0.11430184	−0.10671180	−0.12018854
0	0.02382538	0.01815649
0	0.01702522	0.03152638
	−0.00543948	−0.00662879
	−0.00455690	−0.00257617
		0.00127756
		0.00241187

For any given choice of filters, we may calculate R_a for either the wavelet or scaling functions at each level of the transform. We have chosen to show results for the following combinations of filters:

A (13,19)-tap and (12,16)-tap near-orthogonal odd/even filter sets.
B (13,19)-tap near-orthogonal filters at level 1, 18-tap Q-shift filters at levels ≥ 2.
C (13,19)-tap near-orthogonal filters at level 1, 14-tap Q-shift filters at levels ≥ 2.
D (9,7)-tap Antonini filters at level 1, 18-tap Q-shift filters at levels ≥ 2.
E (9,7)-tap Antonini filters at level 1, 14-tap Q-shift filters at levels ≥ 2.
F (9,7)-tap Antonini filters at level 1, 6-tap Q-shift filters at levels ≥ 2.
G (5,3)-tap LeGall filters at level 1, 6-tap Q-shift filters at levels ≥ 2.

Note that the A filters are for the odd/even dual tree, while filters B to G are for the Q-shift dual tree, in order of decreasing complexity. The C filters are comparable in complexity to the A filters.

For a given dual tree system, R_a may be calculated when the retained coefficients (discussed in Section 4) are from any given level and of either bandpass (wavelet) or lowpass (scaling function) type. Table 3 shows the results of calculating R_a for each level of the DT CWT from 1 to 5 and for both wavelet and scaling function signal paths. It is convenient to represent R_a in dB, using $10 \log_{10} R_a$. Hence a value of −30 dB in the table means that the energy of all the aliased transfer functions is 0.1% of the energy of the main transfer function. The final column of the table (DWT) shows the value of R_a obtained with a conventional fully decimated discrete wavelet transform using just the (13,19)-tap filter set. The approximate complexities of the one-dimensional versions of the dual-tree transforms, relative to the (13,19)-tap DWT, are shown in the second row of the table. These are based only on numbers of non-zero filter taps and not on efficient factorizations.

TABLE 3
Aliasing Energy Ratios, R_a in dB, for Filter Types A to G over Levels 1 to 5

Filters	A	B	C	D	E	F	G	DWT
Complexity	2.0	2.3	2.0	1.9	1.6	1.0	0.7	1.0
Wavelet								
Level 1	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	−9.40
Level 2	−28.25	−31.40	−29.06	−22.96	−21.81	−18.49	−14.11	−3.54
Level 3	−23.62	−27.93	−25.10	−20.32	−18.96	−14.60	−11.00	−3.53
Level 4	−22.96	−31.13	−24.67	−32.08	−24.85	−16.78	−15.80	−3.52
Level 5	−22.81	−31.70	−24.15	−31.88	−24.15	−18.94	−18.77	−3.52
Scaling fn.								
Level 1	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	$-\infty$	−9.40
Level 2	−29.37	−32.50	−30.17	−24.32	−23.19	−19.88	−15.93	−9.38
Level 3	−28.17	−35.88	−29.21	−36.94	−29.33	−21.75	−20.63	−9.37
Level 4	−27.88	−37.14	−28.57	−37.37	−28.56	−24.37	−24.15	−9.37
Level 5	−27.75	−36.00	−28.57	−36.01	−28.57	−24.67	−24.65	−9.37

From the table we see that the longer filters do indeed provide improved shift invariance, and that the complexity of the level 1 filters affects the shift invariance mostly at levels 2 and 3, while the remaining filters affect the shift invariance mostly beyond level 2. There is no aliasing at level 1 for either form of the dual tree, because the two trees behave like a single undecimated tree at this level. The DWT clearly exhibits much worse shift invariance than any of the dual-tree transforms.

The scaling function results tend to be several dB better than the equivalent wavelet results. This is largely due to the lower residual overlap of the lowpass spectra in Fig. 5c compared to the bandpass spectra in Fig. 5d.

The degree of shift invariance of four of the above schemes is illustrated in Fig. 8. In each case, the input is a unit step, shifted to 16 adjacent sampling instants in turn. Each unit step is passed through the forward and inverse version of the chosen wavelet transform. The figure shows the input steps and the components of the inverse transform output signal, reconstructed from the wavelet coefficients at each of levels 1 to 4 in turn and from the scaling function coefficients at level 4. Summing these components reconstructs the input steps perfectly. Good shift invariance is shown when all the 16 output components from a given level are the same shape, independent of shift. It is clear that the DT CWT filters of type B are the best, closely followed by type E. The imperfections of the simple type G filters are fairly obvious in Fig. 8c, while the severe shift dependence of the normal DWT is shown in Fig. 8d. These results follow closely those based on R_a in Table 3.

8. EXTENSION TO *m* DIMENSIONS

Extension of the DT CWT to two dimensions is achieved by separable filtering along columns and then rows. However, if column and row filters both suppress negative frequencies, then only the first quadrant of the 2-D signal spectrum is retained. It is well known from 2-D Fourier transform theory that two adjacent quadrants of the spectrum are

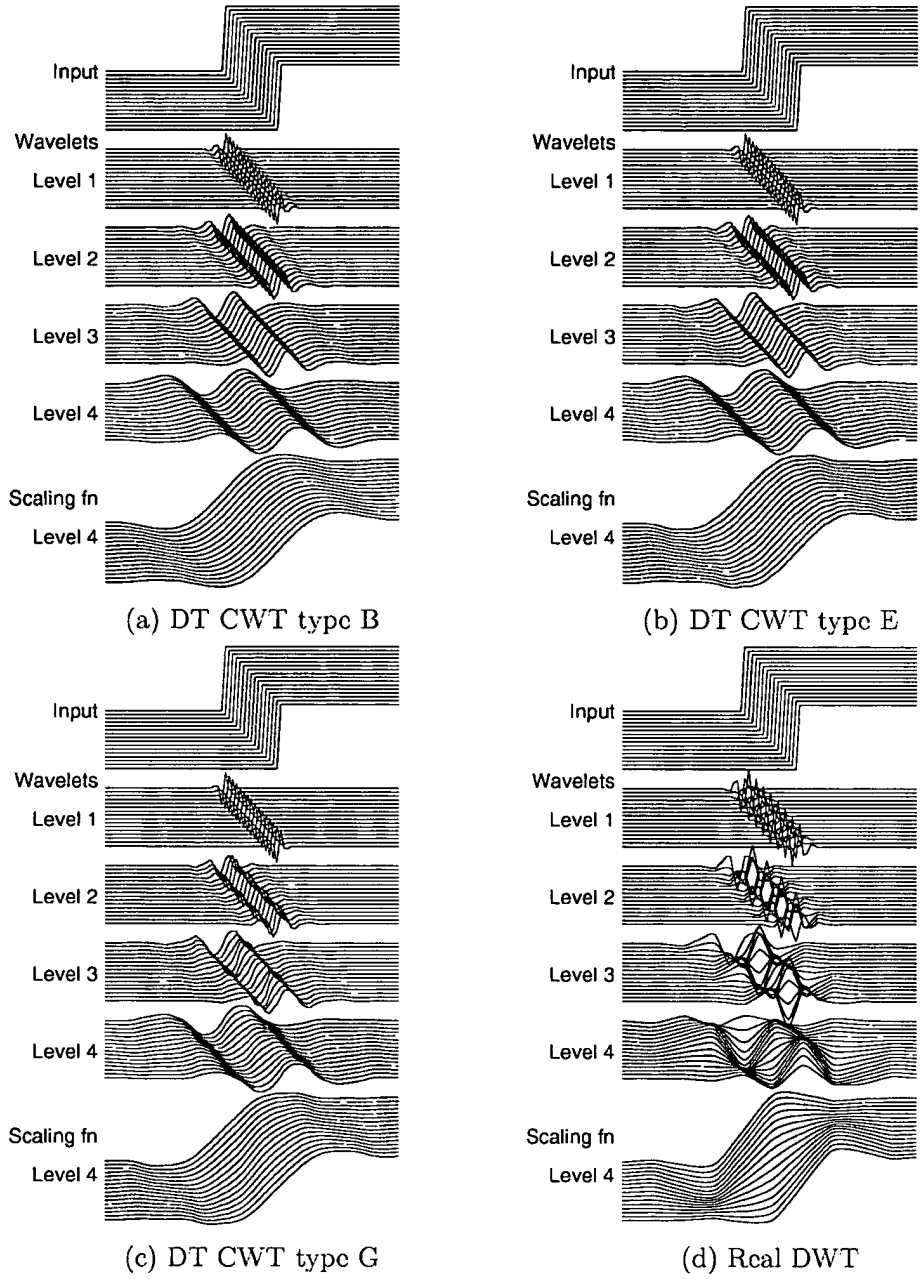

FIG. 8. Wavelet and scaling function components at levels 1 to 4 of 16 shifted step responses of the DT CWT (a, b, and c) and real DWT (d).

required to represent fully a real 2-D signal. Therefore in the DT CWT we also filter with complex conjugates of the row (or column) filters in order to retain a second (or fourth) quadrant of the spectrum. This then gives 4 : 1 redundancy in the transformed 2-D signal. If the signal exists in m dimensions, further conjugate pairs of filters are needed for each additional dimension, leading to a redundancy of 2^m : 1. This process is discussed in more

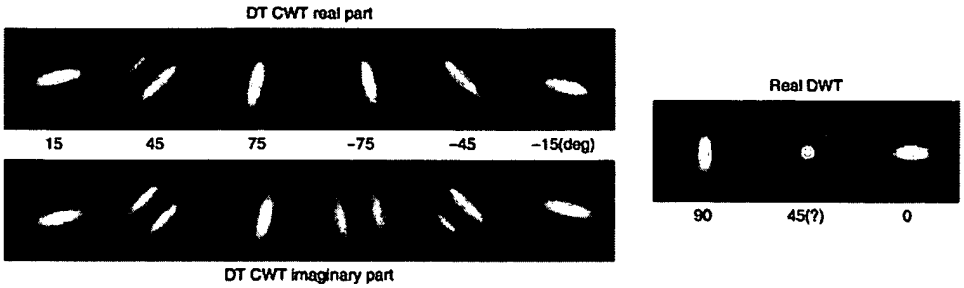

FIG. 9. Impulse responses of 2-D complex wavelet filters (left) and of 2-D real wavelet filters (right), all illustrated at level 4 of the transforms. The complex wavelets provide 6 directionally selective filters, while real wavelets provide 3 filters, only two of which have a dominant direction.

detail in [13]. It is important to note that these redundancy values are *not* affected either by the number of levels of the transform or by whether row and column processing is done at each level before proceeding to the next level (as is usual with the DWT) rather than the other option of doing all levels of row processing first before proceeding to columns, or vice versa.

Complex filters in multiple dimensions can provide true directional selectivity, despite being implemented separably, because they are still able to separate all parts of the m-D frequency space. For example, a 2-D CWT produces six bandpass subimages of complex coefficients at each level, which are strongly oriented at angles of $\pm 15°$, $\pm 45°$, $\pm 75°$, as illustrated by the level 4 impulse responses in Fig. 9. In order to obtain these directional responses, it is necessary to interpret the scaling function (lowpass) coefficients from the two trees as complex pairs (rather than as purely real coefficients at double the rate) so that they can be correctly combined with wavelet (highpass) coefficients from the other dimension, which are also complex, to obtain the filters oriented at $\pm 15°$ and $\pm 75°$. The type C wavelet filters were used in this case.

It is interesting to note that m-dimensional CWTs will produce $(4^m - 2^m)/2$ directional bandpass subbands at each level. In 3-D this gives 28 subbands at each level or scale, which are selective to near-planar surfaces, corresponding to approximately equally spaced points on the surface of a hemisphere.

In Fig. 8, the shift-dependent properties of the DT CWT are compared with the DWT for one-dimensional step functions. In Fig. 10, a similar comparison is made in 2-D, using the DT CWT filters of type C and the same DWT as in Table 3. The input is now an image of a light circular disc on a dark background. The upper row of images, from left to right in Fig. 10, show the components of the output image, reconstructed from the DT CWT wavelet coefficients at levels 1, 2, 3, and 4 and from the scaling function coefficients at level 4. The lower row of images show the equivalent components when the fully decimated DWT is used instead. In the lower row, we see substantial aliasing artifacts, manifested as irregular edges and stripes that are almost normal to the edge of the disc in places. Contrast this with the upper row of DT CWT images, in which artifacts are virtually absent. The smooth and continuous images here demonstrate good shift invariance because all parts of the disc edge are treated equivalently; there is no shift dependence. These images also show good rotational invariance, because each image is using coefficients from all six directional subbands at the given wavelet level. The only rotational dependence is a slight thinning of

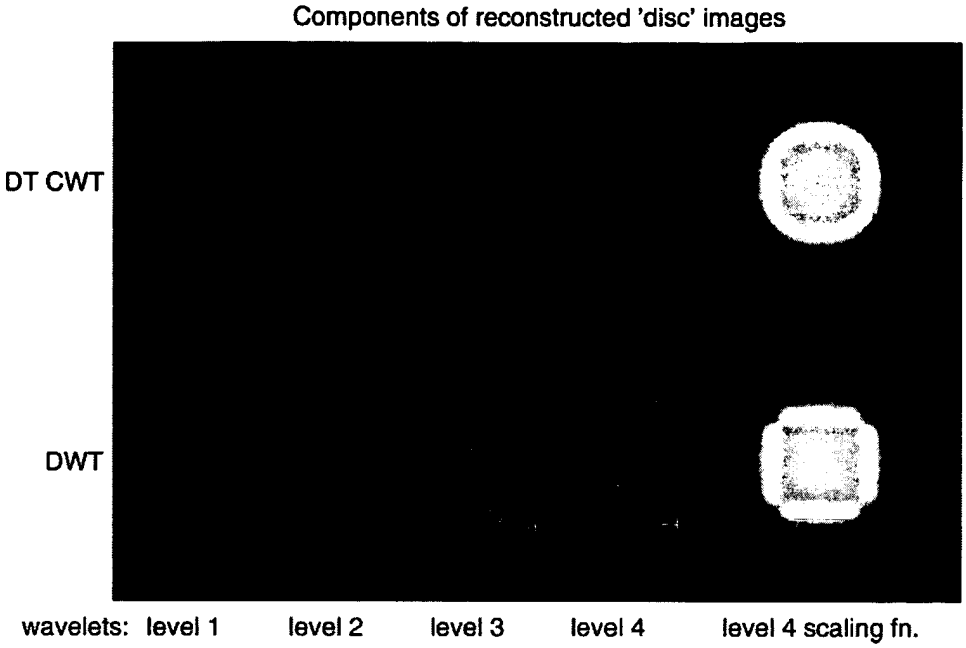

FIG. 10. Components of the reconstructed image of a light circular disc on a dark background, for wavelets and scaling functions at levels 1 to 4, using the 2-D DT CWT (upper row) and 2-D DWT (lower row). Only half of each wavelet image is shown in order to save space. At each wavelet level, all six directional subbands are retained.

the rings of the bandpass images near orientations of $\pm 45°$ and $\pm 135°$, due to the diagonal subbands having higher centre frequencies than the others.

The practical advantages of shift invariant transforms become more obvious if one considers what happens to the reconstructed version of an image, such as the disc in Fig. 10, when some of the wavelet subbands are scaled differently from others. This may be a reduction in gain to achieve denoising or an increase in gain to achieve deblurring. With unity gain for all subbands, summation of either row of images in Fig. 10 produces a perfectly reconstructed disc; but with *unequal* gains, the artifacts of the lower row will tend to appear, while the upper row will just result in uniform blurring or sharpening of edges. Furthermore, by applying gain changes selectively to differently oriented subbands of the CWT, it is possible, for example, to denoise along an edge while sharpening in a direction normal to it.

9. APPLICATIONS

We believe that the shift invariant and directionally selective features of the DT CWT can provide designers with increased flexibility for spatially adaptive filtering of multi-dimensional signals without needing to worry about introduction of undesirable aliasing artifacts. This is likely to be an important feature for many applications, some of which we now consider briefly. In all cases the main signals are images, but may also be video sequences and general 3-D datatsets, such as medical scans and geological seismic data.

Motion estimation and compensation. When two frames of an image sequence are analyzed by the CWT, the phase shifts of the complex coefficients from one frame to the next tend to be proportional to the displacement or motion normal to the orientation of each directional subband. Hence it is possible to determine the displacement vectors at most points in the current frame [7]. Similar principles may also be used to achieve sub-pixel motion compensation, although spline methods may be equally effective in this case.

Denoising and deconvolution. It is well known that wavelet denoising methods such as those of Donoho [9] perform better when implemented with a shift-invariant transform such as the undecimated DWT (UWT). We have shown [11] that the DT CWT performs at least as well as the UWT in this context, and with significantly lower computational cost (3 times that of the basic DWT, compared with $3M$ times for an M-level UWT). In fact the directional properties of the DT CWT allow it to outperform the UWT when the image contains significant diagonal edges. Choi *et al.* [16] have incorporated the DT CWT with a hidden Markov tree (HMT) to achieve further improvements in denoising performance. In this case the HMT could be integrated easily into the DT CWT framework (due to its simple parent-child structure), whereas this is much more difficult for the UWT in which each child has multiple parents. Finally it is straightforward to combine deconvolution with denoising, and Jalobeanu [19] has shown one way of doing this, which involves the use of a *packet* DT CWT at levels 1 and 2 to give improved high frequency selectivity.

Texture analysis and synthesis. The multi-scale and directional properties of the DT CWT make it well suited for texture analysis [14, 15]. In this context it behaves similarly to a multi-scale Gabor filter bank, but is more efficiently implemented. Shift invariance is important too, as it makes the texture feature vectors independent of precise texture location, and Hill [20] has shown how to make the texture features rotationally invariant. Finally the perfect reconstruction features of the DT CWT allow iterative texture synthesis techniques to be used [10].

Segmentation and classification. Using the DT CWT to give combined texture and colour feature vectors at multiple scales allows the development of effective hierarchical segmentation algorithms, which can then be applied to image classification [17, 21]. In [18] Romberg *et al.* have shown how the HMT may be used effectively with the DT CWT in this context, as well as for denoising.

Watermarking. Effective image watermarks must be approximately matched to the local spectral characteristics of the image that is being marked, in order to make the mark as invisible as possible (due to perceptual masking) and as robust as possible against denoising attacks. In [22], we propose a system which achieves this by watermarking in the complex wavelet domain, since it is very easy, using the DT CWT, to provide the spatially varying spectral characteristics that are required for optimal watermarking.

Hence we see that the strengths of the DT CWT stem from its abilities (a) to analyze multi-dimensional signals unambiguously at multiple scales and directions, and (b) to provide spatially adaptive directional filtering which does not suffer from significant aliasing artifacts, all at modest computational cost.

10. CONCLUSIONS

A method of analyzing the shift invariance of the dual-tree complex wavelet transform has been presented, based on the aliasing energy ratio R_a. The DT CWT is shown to possess good shift invariance properties, given suitably designed biorthogonal or orthogonal wavelet filters. These properties extend to multiple dimensions, where either approximate rotational invariance or good directional selectivity can also be provided. The computational advantages of the complex wavelet approach over undecimated wavelets are particularly significant with multi-dimensional signals. The filter design problem is also discussed and the coefficients of suitable sets of filters for the DT CWT are provided.

A range of applications of the dual tree transform has been summarized and a bibliography has been provided to papers describing these applications, a number of which may be viewed at the author's home page: www-sigproc.eng.cam.ac.uk/~ngk.

REFERENCES

1. S. G. Mallat, A theory for multiresolution signal decomposition: The wavelet representation, *IEEE Trans. PAMI* **11** (1989), 674–693.
2. S. G. Mallat, "A Wavelet Tour of Signal Processing," Academic Press, San Diego, 1998.
3. P. Burt and E. H. Adelson, The Laplacian pyramid as a compact image code, *IEEE Trans. Comm.* **31** (1983), 532–540.
4. E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, Shiftable multiscale transforms, *IEEE Trans. Inform. Theory* **38** (1992), 587–607.
5. R. H. Bamberger and M. J. T. Smith, A filter bank for the directional decomposition of images: Theory and design, *IEEE Trans. Signal Process.* **40** (1992), 882–893.
6. P. P. Vaidyanathan and P.-Q. Hoang, Lattice structures for optimal design and robust implementation of two-channel perfect reconstruction QMF banks, *IEEE Trans. ASSP* (1988), 81–94.
7. J. F. A. Magarey and N. G. Kingsbury, Motion estimation using a complex-valued wavelet transform, *IEEE Trans. Signal Process.* **46** (1998).
8. D. B. H. Tay and N. G. Kingsbury, Flexible design of multidimensional perfect reconstruction FIR 2-band filters using transformations of variables, *IEEE Trans. Image Process.* **2** (1993), 466–480.
9. D. L. Donoho, De-noising by soft thresholding, *IEEE Trans. Inform. Theory* **41** (1995), 613–627.
10. N. G. Kingsbury, The dual-tree complex wavelet transform: A new technique for shift invariance and directional filters, *in* "Proc. 8th IEEE DSP Workshop, Bryce Canyon, August 1998."
11. N. G. Kingsbury, The dual-tree complex wavelet transform: A new efficient tool for image restoration and enhancement, *in* "Proc. EUSIPCO 98, Rhodes, September 1998."
12. N. G. Kingsbury, Shift invariant properties of the dual-tree complex wavelet transform, *in* "Proc. ICASSP 99, Phoenix, AZ, paper SPTM 3.6, March 16–19, 1999."
13. N. G. Kingsbury, Image processing with complex wavelets, *Philos. Trans. Roy. Soc. London Ser. A* **357** (1999), 2543–2560.
14. S. Hatipoglu, S. K. Mitra, and N. G. Kingsbury, Texture classification using dual-tree complex wavelet transform, *in* "Proc. 7th International IEE Conference on Image Processing and Its Applications, Manchester, England, July 12–15, 1999," pp. 344–347.
15. P. F. C. de Rivaz and N. G. Kingsbury, Complex wavelet features for fast texture image retrieval, *in* "Proc. IEEE Conf. on Image Proc., Kobe, Japan, October 25–28, 1999."
16. H. Choi, J. Romberg, R. Baraniuk, and N. G. Kingsbury, Hidden Markov tree modelling of complex wavelet transforms, *in* "Proc. ICASSP 2000, Instanbul, June 6–9, 2000."
17. A. H. Kam, T. T. Ng, N. G. Kingsbury, and W. J. Fitzgerald, Content based image retrieval through object extraction and querying, *in* "Proc. IEEE Workshop on Content-based Access of Image and Video Libraries (CBAIVL-2000), Hilton Head, South Carolina, June 12, 2000."

18. J. Romberg, H. Choi, R. Baraniuk, and N. G. Kingsbury, Multiscale classification using complex wavelets, *in* "Proc. ICIP 2000, Vancouver, September 2000."
19. A. Jalobeanu, L. Blanc-Feraud, and J. Zerubia, Satellite image deconvolution using complex wavelet packets, *in* "Proc. ICIP 2000, Vancouver, September 2000."
20. P. Hill and D. Bull, Rotationally invariant texture features using the dual-tree complex wavelet transform, *in* "Proc. ICIP 2000, Vancouver, September 2000."
21. P. F. C. de Rivaz and N. G. Kingsbury, Fast segmentation using level set curves of complex wavelet surfaces, *in* "Proc. ICIP 2000, Vancouver, September 2000."
22. P. Loo and N. G. Kingsbury, Digital watermarking using complex wavelets, *in* "Proc. ICIP 2000, Vancouver, September 2000."